



AFRL-AFOSR-VA-TR-2015-0393

A Distributed Representation of Remembered Time

Marc Howard
TRUSTEES OF BOSTON UNIVERSITY

11/19/2015
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE		Form Approved OMB No. 0704-0188
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Executive Services, Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.</p>		
1. REPORT DATE (DD-MM-YYYY) 23-11-2015	2. REPORT TYPE Final Performance	3. DATES COVERED (From - To) 01-07-2012 to 30-06-2015
4. TITLE AND SUBTITLE A Distributed Representation of Remembered Time		5a. CONTRACT NUMBER 5b. GRANT NUMBER FA9550-12-1-0369
		5c. PROGRAM ELEMENT NUMBER 61102F
6. AUTHOR(S) Marc Howard		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) TRUSTEES OF BOSTON UNIVERSITY 1 SILBER WAY BOSTON, MA 02215-1390 US		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AF Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203		10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR RTA2
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT A DISTRIBUTION UNLIMITED: PB Public Release		
13. SUPPLEMENTARY NOTES		
14. ABSTRACT <p>The goal of the award was to extend a mathematical framework for representing time to also construct scale-invariant representations of space. The extensive neurophysiological work done on the hippocampal place code (which resulted in a Nobel Prize in 2014) places physical constraints on this mechanism. We accomplished this goal by developing a computational framework that describes a wide range of functional cellular correlates in the hippocampus and related brain regions. Critically, this cellular-level model is connected to behavioral memory performance via cognitive models that take in the mathematical form of the representation. We have exceeded the initial goals of the project by developing a representation of numerosity as well as space using the same framework and a detailed model for function translation that is constrained by neurophysiological data from the hippocampal theta oscillation literature. This is an important development, as it opens up new frontiers in cognitive computation. Taken together, the work performed in this three year period was crucial in developing a mathematical model for cognitive operations that is constrained by neurophysiology. This model connects data from the subcellular level, to the systems neuroscience level to cognitive modeling</p>		
15. SUBJECT TERMS Time		

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

DISTRIBUTION A: Distribution approved for public release.

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Marc Howard
Unclassified	Unclassified	Unclassified	UU		19b. TELEPHONE NUMBER (Include area code) 315-443-1864

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

DISTRIBUTION A: Distribution approved for public release.

A Distributed Representation of Remembered Time: Final report

Marc W. Howard

Department of Psychological and Brain Sciences

Center for Memory and Brain

Boston University

FA9550-12-1-0369

Abstract

The goal of the award was to extend a mathematical framework for representing time to also construct scale-invariant representations of space. The extensive neurophysiological work done on the hippocampal place code (which resulted in a Nobel Prize in 2014) places physical constraints on this mechanism. We accomplished this goal by developing a computational framework that describes a wide range of functional cellular correlates in the hippocampus and related brain regions. Critically, this cellular-level model is connected to behavioral memory performance *via* cognitive models that take in the mathematical form of the representation. We have exceeded the initial goals of the project by developing a representation of numerosity as well as space using the same framework and a detailed model for function translation that is constrained by neurophysiological data from the hippocampal theta oscillation literature. This is an important development, as it opens up new frontiers in cognitive computation. Taken together, the work performed in this three year period was crucial in developing a mathematical model for cognitive operations that is constrained by neurophysiology. This model connects data from the subcellular level, to the systems neuroscience level to cognitive modeling. The mathematical framework is capable of forming the basis of a general brain-inspired cognitive computer.

The award supported a diverse range of theoretical work leading towards an integrative physical theory of brain function ranging from subcellular processes to systems neurophysiology to cognition and behavior. In order to disseminate these results broadly, we have published papers in a wide range of disciplines, including outlets that specialize in machine learning and artificial intelligence (Shankar & Howard, 2013; Howard, Shankar, & Tiganj, 2015; Shankar, 2015), cognitive science (Howard & Eichenbaum, 2015, 2013; Shankar, 2014), neurobiology (Howard et al., 2014; Tiganj, Hasselmo, & Howard, 2015;

Howard & Eichenbaum, 2015), and even theoretical physics (Shankar, Singh, & Howard, 2015). This document provides a narrative description of the state of the theoretical development at the time of this writing and briefly describes future directions we hope to pursue.

Theoretical overview

We have developed a formal mechanism that naturally computes a temporal memory that can be self-sufficiently evolved. The key insight is that while the timeline itself cannot be evolved self-sufficiently in time, the Laplace transform of the timeline can be (Shankar & Howard, 2013). The model can be understood as a two-layer feedforward architecture (Fig 1a). At each moment τ a single input node f projects to a set of units $F(s)$ that store the Laplace transform up to the current moment; s indexes the different units. Through a local set of feed forward connections (represented by an operator \mathbf{L}_k^{-1}), the second layer approximately inverts the encoded Laplace transform to represent a fuzzy reconstruction of the actual stimulus history itself, $\tilde{f}^*(\tau)$.

There are two equations necessary to construct the representation. First, we must encode and update the Laplace transform, with real coefficients s . This can be accomplished if each unit in $F(s)$ is updated according to

$$\frac{dF(s, \tau)}{d\tau} = \alpha(\tau) [-sF(s, \tau) + f(\tau)], \quad (1)$$

where s indexes the unit. Note that this equation is both time- and space-local in that each unit only needs to know its own value (local in space) at the previous time step (local in time) in order to compute its new value. This ability to update self-sufficiently decouples each value of s from its neighbors and enables the distribution of units across s to implement Weber-Fechner coding (Shankar & Howard, 2013; Shankar, 2015). The factor $\alpha(\tau)$ modulates all of the units in $F(s)$ with some value that varies with time. Assigning different functions $\alpha(\tau)$ enables the same coding scheme to describe many different quantities. If $\alpha(\tau)$ is constant, $F(s)$ codes for the Laplace transform (with real positive coefficients) of f as a function of time. If $\alpha(\tau)$ is set to velocity in a particular direction, $dx/d\tau$, then $F(s)$ codes for the Laplace transform of $f(x)$. By choosing $\alpha(\tau)$ appropriately, this framework can be used to describe functional correlates of a wide range of time cells and place cells in the hippocampus (Howard et al., 2014; Howard & Eichenbaum, 2015). As described later, a detailed biophysical simulation (Tiganj et al., 2015) shows that Equation 1 could be implemented using a long-lasting current known to support persistent firing cells in slice preparations of entorhinal and perirhinal cortex (Egorov, Hamam, Fransén, Hasselmo, & Alonso, 2002; Fransén, Tahvildari, Egorov, Hasselmo, & Alonso, 2006; Navaroli, Zhao, Boguszewski, & Brown, 2011).

Given the Laplace transform, it remains to approximate the inversion of the transform. The Post inversion formula (Post, 1930) provides a neurally-plausible way to invert the transform:

$$\tilde{f}^*(\tau, \tau) \equiv \mathbf{L}_k^{-1} F(s, \tau) \equiv C_k s^{k+1} \frac{d}{ds^k}, \quad (2)$$

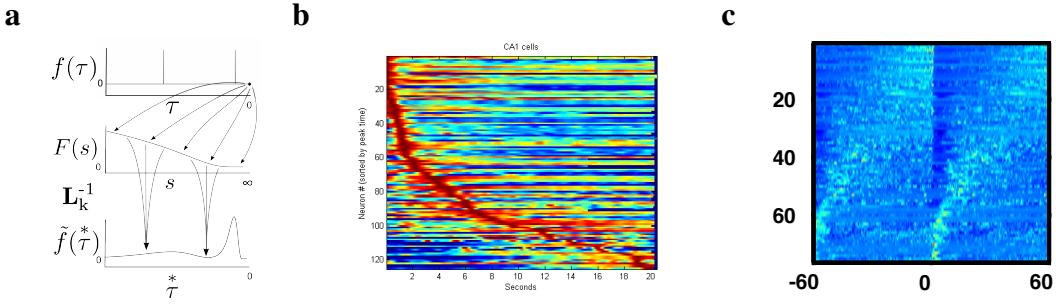


Figure 1. Temporal representation in the brain. **a.** Schematic of the model for encoding a temporal history $f(\tau)$. At each time step, input from a single node provides input to a set of nodes $F(s)$. Each node of F is indexed by a different value of s which can be identified with the real Laplace variable. Nodes in $F(s)$ project locally to another set of nodes in $\tilde{f}(\tau^*)$ via an operator L_k^{-1} . The nodes in \tilde{f} approximate the original function $f(\tau)$. The error in $\tilde{f}(\tau^*)$ is scale invariant. We choose the distribution of nodes across s and thus also τ^* to implement Weber-Fechner spacing (not shown). **b.** Time cells in the hippocampus. Firing rate for each cell is shown as a function of time during the delay period of a memory task. Each row is the average firing rate of an individual cell as a function of time; the cells are sorted on their time of peak firing. If the width of receptive fields and number density was constant, this plot would appear as a straight ridge with constant width. The curvature and spread are qualitatively consistent with Weber-Fechner scaling and statistically reliable across a range of time cell experiments. Unpublished data from the Eichenbaum lab. **c.** Time cells in the striatum show the same qualitative properties. Time (on the x axis) is in seconds. Firing rate is shown prior to (left) and after (right) food delivery in a fixed interval operant task. Food delivery was approximately periodic so there was another stimulus at approximately -60. Here the color scheme is different and the order of sorting is different, but the same qualitative properties hold. From Mello, et al., (2015).

where k is an integer parameter that controls the accuracy of the inversion and C_k is just a normalization constant that depends on k . The operator L_k^{-1} is not only linear but also local—each value of τ^* is in a one-to-one relationship with s (more precisely each value of τ^* is proportional to the time constant determined by s). The time-dependence of \tilde{f} is attributable to the derivative operator, which can be approximated with on the order of k near neighbors around s . Fortunately, L_k^{-1} is both neurally plausible— $k = 2$ can be implemented with a Mexican hat feedforward projection—and computationally advantageous, giving rise to scale-invariant representation such that the width of temporal receptive fields proportional to τ^* . Because both Eqs. 1 and 2 are local, this gives us freedom to choose the distribution of s . It can be shown that Weber-Fechner spacing, such that the difference between adjacent values of τ^* goes up linearly with τ^* , is optimal for long-range correlated signals (Shankar & Howard, 2013).

Cognitive models for learning and memory

When $\alpha(\tau)$ is constant over time, Eqs. 1 and 2 generate a representation as a function of time. Howard, Shankar, Aue, and Criss (2015) developed relatively detailed cognitive models of a number of memory tasks using this scale-invariant time-line. Much of contemporary memory research focuses on different “forms” of memory. We demonstrated that if provided with a Weber-Fechner spaced representation of temporal history, different forms of memory could be understood as different operations on the same underlying representation. For instance, classic short-term working memory tasks (Hacker, 1980) can be understood as the result of scanning along an internal timeline. In contrast, contiguity effects in episodic memory tasks (Kahana, Howard, & Polyn, 2008; Howard, Youker, & Venkatadass, 2008) can be understood as the result of a “jump back in time” in which a remembered stimulus recovers the state of the history when that stimulus was previously experienced (Howard, Viskontas, Shankar, & Fried, 2012; Manning, Polyn, Litt, Baltuch, & Kahana, 2011). In addition to quantitative descriptions of long- and short-term memory tasks in humans, the Howard, Shankar, Aue, and Criss (2015) paper also provided an account of animal conditioning tasks (Cole, Barnet, & Miller, 1995) and a review of the neurophysiological data on time cells available at that time.

Representation of time and space in the hippocampus

Howard et al. (2014) generalized the approach for constructing a scale-invariant timeline to generating representations of spatial location and compared the model’s predictions to neurophysiological findings from the rodent hippocampus and related regions (Fig. 2, Howard et al., 2014 see also Howard & Eichenbaum, 2015). Choosing different functions for $\alpha(\tau)$ and different inputs enables us to reuse the same coding scheme for time, spatial position, ordinal position, or numerosity.

When $\alpha(\tau)$ is constant over time, \tilde{f} can be understood as a scale-invariant timeline. If hippocampal time cells are the neural mechanism constructing that representation, then this makes several quantitative predictions that can be tested. One of the contributions of the (Howard et al., 2014) paper was to test these predictions. First, the width of time fields should grow with the cell’s time of peak firing. Second, time fields should be asymmetric. We showed that real neurons, reported initially by our collaborators MacDonald, Lepage, Eden, and Eichenbaum (2011), exhibit both of these properties. We did not find evidence for another prediction, that the number of cells coding for distant times is smaller than the number of cells coding for more recent times, but subsequent recordings have shown this property in the hippocampus and striatum (Fig. 1).

In addition to testing predictions about time cells, we also reported simulations of a range of place cell phenomena. For these simulations, we took actual paths navigated by rats in various environments and set $f(\tau)$ and $\alpha(\tau)$ appropriately. If we set $\alpha(\tau) = dx/d\tau$, $\tilde{f}(x^*)$ approximates a function of x . If $\alpha(\tau)$ is set to the (one-dimensional) velocity of the animal in an open environment, and if $f(\tau)$ is set to be non-zero when the animal encounters a specific boundary of the enclosure, the model predicts one-dimensional place fields

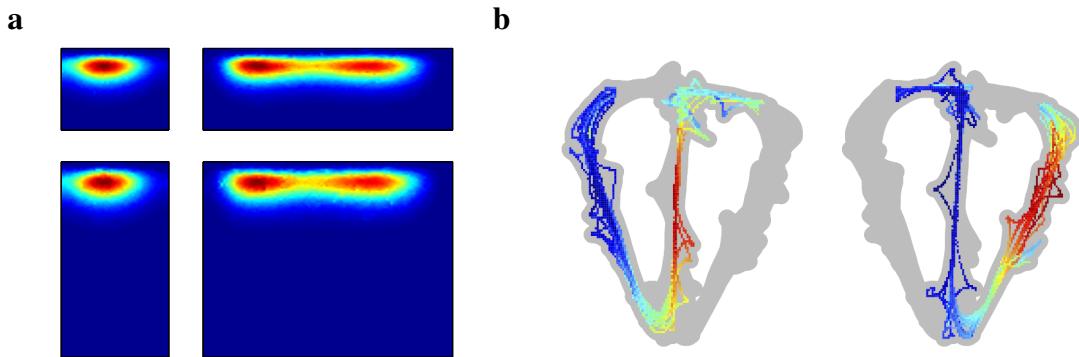


Figure 2. Model for time cells can also generalize to place cells. Both panels from Howard, et al., (2014). **a.** Simulated place cell generated from the model. The four panels show firing rate maps for a place cell in four deformations of a recording enclosure. The simulated cell is constructed as a non-linear combination of the output of two one-dimensional “boundary vector cells.” One of the boundary vector cells codes for distance from the north wall. The other codes for distance from the Eastern wall. This resembles experimental findings from Burgess & O’Keefe (1996). **b.** The model accounts for path-dependent firing, such as shown by “splitter cells.” The animal runs along the maze in a “figure 8” pattern so that the movements along the center arm can be preceded by a visit to either of the other arms (left vs right). The firing rate of a simulated cell are shown in false color. Note that the cell fires along the center arm when coming from the West arm (left) but not when coming from the East arm (right). Paths were provided from an experiment by Lipton, et al., (2007). This resembles experimental findings by Wood et al., (2000).

that appear as “strips” through a 2-D environment. This property closely resembles those described for boundary vector cells, which have been observed in the subiculum (Lever, Burton, Jeewajee, O’Keefe, & Burgess, 2009) and have long been hypothesized to provide inputs to classic hippocampal place cells (O’Keefe & Burgess, 1996; Hartley, Burgess, Lever, Cacucci, & O’Keefe, 2000; Barry et al., 2006). Not surprisingly, by taking conjunctions of our simulated boundary vector cells, we were able to provide firing correlates that resemble many classical findings for hippocampal place fields (Fig. 2a). In this way, the Laplace transform modulated by velocity becomes a mechanism to answer the question of how allocentric position is computed in the hippocampal formation.

In addition to time cells and place cells, the same computational model can describe more elaborate firing correlates as well (see also Howard & Eichenbaum, 2015). Many “place cells” in the hippocampus do not fire only in dependence on position, but in a way that depends on the history leading up to that location. For instance Figure 2b shows a simulation of “splitter cells.” In this experiment animals run alternately turn left and then right at the end of the central arm. This means the animal may run along the central arm having come from the left or having come from the right. Splitter cells fire at a particular location on one type of path or the other. The instantaneous position (and velocity) is the same, but the firing of splitter cells differentiates location based on history. In order to simulate this, we set $f(\tau)$ to the animal’s current head direction and set $\alpha(\tau)$ to the current speed. Note

that the simulated cell fires on the central arm on only one type of journey, resembling the canonical splitter cell phenomenon (see also Frank, Brown, & Wilson, 2000). In this way, the set of all cells in \tilde{f} code for a scale-invariant ordinal representation of the sequence of movements along the path. Of course ordinal representations can also be understood as representations of numerosity.

Detailed biophysical implementation of Eq. 1

The mathematical framework described thus far represents something of a challenge to the conventional understanding of neural dynamics. In order to implement Eq. 1 to describe time cells that fire at latencies of a minute or more, we must be able to generate functional time constants that are much longer than membrane time constants. One possibility is that the “units” in $F(s)$ are not single neurons, but rather populations of recurrently connected neurons with eigenvalue less than 1. Moreover, in order to build representations of space and number, these time constants must be able to be rapidly modulated by a continuously-varying signal $\alpha(\tau)$. Inspired by these constraints, Tiganj et al. (2015) instead pursued the hypothesis that units in $F(s)$ correspond roughly to single neurons. If single neurons can implement Eq. 1, this vastly expands the capacity of the network relative to a population of neurons with recurrent connections. It also connects to a large body of work on persistently-firing neurons from the slice literature.

Persistent firing neurons have effectively infinite functional time constants (Fig. 3a; Egorov et al., 2002). This sustained firing in slice preparations is a consequence of a nonspecific calcium-sensitive cation (CAN) current that depends on the concentration of acetylcholine in the bath. Neurons with similar properties have been observed in medial entorhinal cortex (Egorov et al., 2002), amygdala (Egorov, Unsicker, & von Bohlen und Halbach, 2006) and perirhinal cortex (Navaroli et al., 2011), suggesting that persistent firing cells are an important component of the computational function of the MTL. A prominent model of persistent firing cells (Fransén et al., 2006) hypothesizes that stable persistent firing is regulated by a metabolic cascade that maintains stable calcium concentration in the absence of a large perturbation. Equation 1 does not require stable persistent firing, but rather gradually decaying firing. The basic idea of the computational model of (Tiganj et al., 2015) is that if calcium concentration responded to smaller perturbations, this would enable decaying persistent firing. In this model, in the absence of other processes, intracellular calcium decays exponentially with some relatively short time constant. However, spikes cause the entry of calcium. The CAN current is sensitive to the presence of calcium and causes more spiking. In this way, spiking causes the actual calcium concentration, and spiking, to decay much more slowly than the intrinsic time constant of calcium clearance. Using parameter settings chosen either from empirical values or those chosen from other computational models, Tiganj et al. (2015) showed that this mechanism enables the model to account for exponentially-decaying firing with arbitrarily large time constant (Fig. 3b).

The derivation of Tiganj et al. (2015) identified several factors that can alter the functional time constant. Intriguingly, that paper showed that, in the presence of the CAN current, any manipulation that changes the slope of the f-I curve (relating firing rate to the

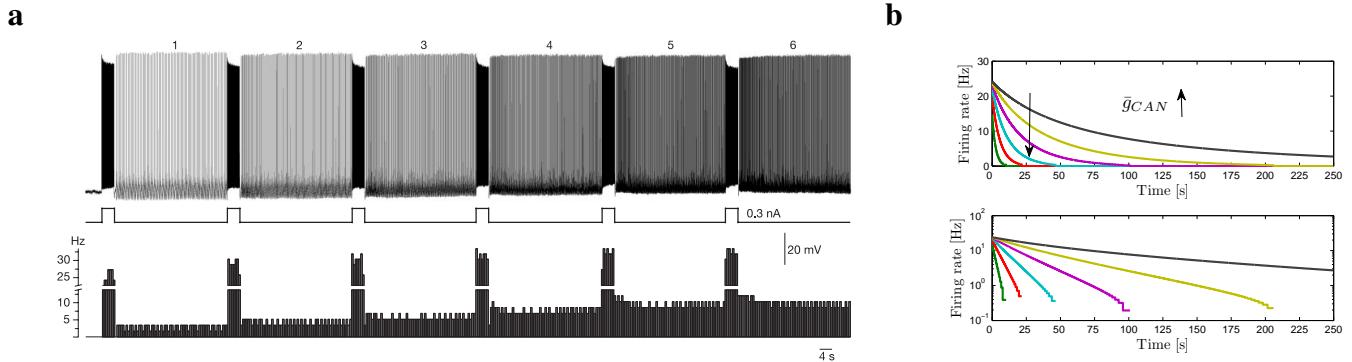


Figure 3. A simple biophysically plausible mechanism for generating long, modifiable functional time constants. **a.** Persistent firing cells observed in slice preparations in entorhinal cortex, perirhinal cortex, and amygdala have an effectively infinite functional time constants. The top trace shows voltage across the membrane of a synaptically-isolated cell. As current is injected, the cell starts firing at a stable firing rate. As additional current injections are provided, the cell adapts to a new, stable firing rate (bottom). Note that the scale bar is 4 s. This persistent activity depends on a nonspecific cation (CAN) current which is modulated by calcium. After Egorov, et al., (2002). **b.** A biophysically-plausible simulation can adapt the effectively infinite time constants to a range of finite values that can be modulated by several factors. Here, firing rate from an integrate-and-fire simulation is shown as a function of time for several values of the conductance of the CAN current. The bottom plot shows log of the firing rate; the straight lines indicate that the firing is exponential but with different time constants.

magnitude of the CAN current) has the effect of modifying the time constant, implementing the variation necessary for $\alpha(\tau)$. There is extensive neurophysiological evidence that such divisive inhibition can be induced by external inputs (e.g., Chance, Abbott, & Reyes, 2002; Longtin, Doiron, & Bulsara, 2002 for a review, see Silver, 2010).

Translation of 1-D spatial and temporal representations via hippocampal theta oscillations

Access to the Laplace transform of f enables efficient neural computation. Shankar et al. (2015) described a mathematical and neurophysiological model for translation of functions $\tilde{f}(\tau, \tau + \delta)$ to construct an estimate of the future (Fig. 4a; Shankar et al., 2015). Although the model applies equally well to time and one-dimensional position, we will restrict our attention to time in this discussion for simplicity. The model assumes that at each moment stimuli are associated to the state of the timeline at the time they are experienced. Suppose we want to predict what stimulus will occur 10 s into the future. If we had access to the future state of the timeline, we could determine what will happen at that future time using simple association. The method for estimating the future proceeds by translating the current state of the timeline into the past (i.e., in the future the present will be in the past). In the Laplace domain, translation is mathematically straightforward and can be understood as manipulating the weights in L_k^{-1} . Critically, the computational time

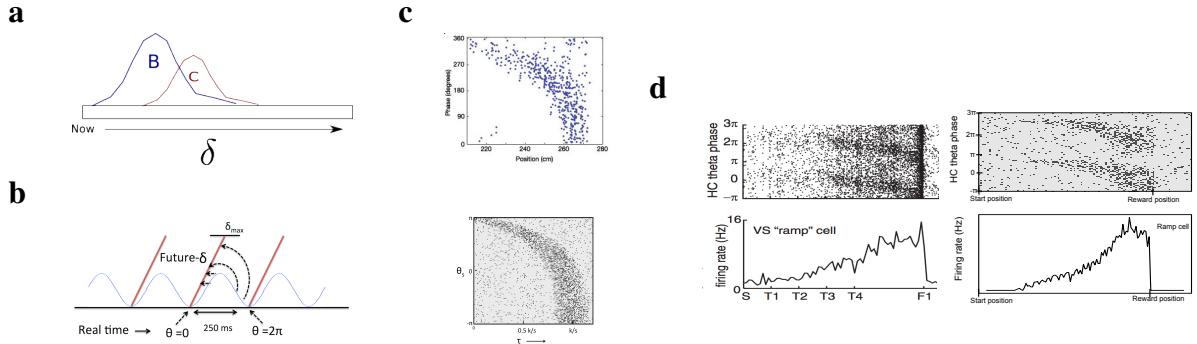


Figure 4. A neural mechanism for constructing an ordered representation of future time comports with phase precession data. **a.** The goal of translation is to generate an estimated timeline of future events. Here δ indexes the distance into the future; following presentation of A, B is predicted earlier in the future than C. **b.** The mapping hypothesis is that δ is swept from zero to some large value within each theta cycle. The rate at which δ changes within a theta cycle accelerates as the cycle proceeds. **c.** Phase precession data from a place cell in the rodent hippocampus (top) and in the model (bottom). Each spike fired by a particular neuron is an individual point. The animal's location within the place field as it traverses left-to-right (x axis) has an effect on the phase of firing relative to theta (y axis). Representative data from Mehta, et al., (2002). **d.** Theta phase precession shows properties resembling translation to different future points of the trajectory within a theta cycle. Top: neurophysiological data from Kjelstrup et al., (2008). Place cells from different positions along the dorsoventral axis of the hippocampus have place cells of different size. However, cells at all scales still precess over the same range of phases. Bottom: model predictions from Shankar, et al., (2015) show the same qualitative patterns. **d.** A phase precessing cell in the rodent ventral striatum as the animal traverses a multiple-choice maze. Left, data from van der Meer et al. (2011). At F1, a reward was presented. The start of the maze is labeled S. The cell ramps up its firing over the entire maze, phase precessing relative to hippocampal theta. Right, mathematical model. In this case, the simulated cell corresponds to prediction of the reward, which gradually ramps up over the length of the maze.

it takes to translate a certain distance does not depend linearly on that distance, enabling logarithmic compression of the future.

The mapping hypothesis is that translation takes place within each theta cycle (Fig. 4b). We derived the properties of translation constrained by a set of findings from the place cell literature operationalized as phenomenological equations (O'Keefe & Recce, 1993; Lubenov & Siapas, 2009; Kjelstrup et al., 2008; van der Meer & Redish, 2011). In particular, we used the constraints that the scale of place cells varies along the dorsoventral axis of the hippocampus (Jung, Wiener, & McNaughton, 1994; Kjelstrup et al., 2008), and that theta is a traveling wave along the dorsoventral axis of the hippocampus (Lubenov & Siapas, 2009; Patel, Fujisawa, Berényi, Royer, & Buzsáki, 2012). Given these constraints, the derivation requires that the representation of the predicted future must be organized on a logarithmically-compressed axis, closely analogous to the logarithmically-compressed

representation of past time (Shankar & Howard, 2013; Howard, Shankar, Aue, & Criss, 2015). It naturally accounts for the finding that hippocampal phase precession proceeds across the same range of phases—and thus larger spatial scales—across the dorsoventral axis (Kjelstrup et al., 2008). In addition, cells coding predicted future outcomes should show ramping activity from the time the cue becomes available to the time the outcome is realized, phase precessing throughout the duration. This prediction is qualitatively consistent with findings from ramping cells in the ventral striatum (Fig. 4d; van der Meer & Redish, 2011).

Future Directions: General Cognitive Computing

Howard, Shankar, and Tiganj (2015) sketched a research strategy for developing this mathematical approach into a general neural computer for one-dimensional quantities. Cognitive computation in the brain is fast, efficient and flexible. Emulating this ability would result in extremely important technological advances. A general computational framework should be able to operate on a wide range of content without learning each exemplar. Such a framework should generalize across not only different specific operands but also across sensory domains, providing a general computational language for cortical computation. Mathematical operations are an important aspect of symbolic processing. Because of the combinatorics of these problems, learning each set of operands and the appropriate outcome is not feasible. This concern, along with the observation that cognition exhibits properties such as productivity and compositionality that are difficult to imagine in standard connectionist architectures has long been seen as an intrinsic limitation on neural computing (e.g., Fodor & Pylyshyn, 1988).

If time, space, and number, as well as sensory representations share a common coding scheme, then mechanisms for computing with representations of that form could be reused across many types of information. Translation is an example of a general computation that can be used on representations from any number of modalities. Translation of functions of time can be used to anticipate the future to inform decision-making in the present; translation of functions of other variables can be used to imagine alternative states of the world to inform decision-making in the world in its current state.

In the same way that point-wise multiplications in the Laplace domain can implement translation of any function, (Howard, Shankar, & Tiganj, 2015) also showed that it is possible to perform *addition* and *subtraction* operations on any two functions by point-wise parallel computations in the Laplace domain. If we have two functions, $f(x)$ and $g(x)$, that represent probability density functions over x , then the sum of those two density functions can be represented by the convolution of f and g . Whereas convolution in x is an $O(N^2)$ operation (where N is the number of cells supporting x), it turns out that the convolution of f and g can be computed using pairwise multiplication in the Laplace domain. In particular, the Laplace transform of the convolution is just $F(s)G(s)$, which requires $O(N)$ operations which can be performed in parallel. The inverse operator \mathbf{L}_k^{-1} can be used to invert the transform, recovering an approximation of the convolution. Again, because \mathbf{L}_k^{-1} is linear, this can be performed in parallel. The subtraction of f and g can be computed

using cross-correlation, which is similarly efficient in the Laplace domain.

We have not yet carefully considered the neurophysiological substrates that could support these arithmetic operations. However, the computational efficiency of performing these operations in the Laplace domain is considerable. Given these considerations, it may be reasonable for the brain to encode the Laplace transform even for variables that are provided along a Weber-Fechner scale due to the property of the sensory receptors.

This suggests a general strategy for implementing symbolic computation in neural architectures. It is clear that arithmetic obeys the properties of compositionality and productivity (modulo edge effects). If the result of an addition operation is a function with the same neural code as the addends, then one can in principle represent an effectively infinite number of possible problems. For instance, given only two input functions f and g one could compute $f + g$, or $(f + g) + g$, or $(f + f) + g$, etc.

Given a set of simple operations one one-dimensional representations of functions such as time, space, and number, a natural extension is to develop an N -dimensional framework with similar flexibility. This leads to the question of how the brain represents higher-dimensional quantities. We have seen that one-dimensional quantities x are in many cases represented by supporting different parts of x with the firing rate of different neurons. How about functions over two-dimensional quantities? Take for example the hippocampal place code, which represents 2-D (x, y) position within an enclosure. The brain might have chosen to have a set of N cells that support position along the x direction and another N cells that support y position. If that were the case, we would require $2N$ cells, each of which would show place fields that are “stripes” through the environment. In contrast, the real place code is constructed by cells that approximate radial basis functions, firing in some circumscribed region of the environment. This coding scheme requires N^2 cells to tile the surface, but is nonetheless preferred. Of course, general cognitive computation, for instance operating over a semantic space, would require many more than two dimensions. Consideration of neural coding in the prefrontal cortex (Rigotti et al., 2013) and hippocampus (Komorowski et al., 2013) suggests that the brain supports multiple dimensions with the same sort of coding scheme, at least in these “high level” cognitive processing centers. A set of simple operations that work on high-dimensional representations, including semantic representations, could be used to develop a general neurally-inspired cognitive computer.

Publications Supported by the Award

Howard, M. W., & Eichenbaum, H. (2013). The hippocampus, time, and memory across scales. *Journal of Experimental Psychology: General*, 142(4), 1211-30. doi: 10.1037/a0033621

Howard, M. W., & Eichenbaum, H. (2015). Time and space in the hippocampus. *Brain Research*, 1621, 345-354.

Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *Journal of Neuroscience*, 34(13), 4692-707. doi: 10.1523/JNEUROSCI.5808-12.2014

Howard, M. W., Shankar, K. H., Aue, W., & Criss, A. H. (2015). A distributed representation of internal time. *Psychological Review*, 122(1), 24-53.

Howard, M. W., Shankar, K. H., & Tiganj, Z. (2015). Efficient neural computation in the Laplace domain. In *Cognitive computation: Integrating neural and symbolic computation*.

Shankar, K. H. (2014). Quantum random walks and decision making. *Topics in cognitive science*, 6(1), 108–113.

Shankar, K. H. (2015). Generic construction of scale-invariantly coarse grained memory. *Lecture Notes in Artificial intelligence*, 8955, 175-184.

Shankar, K. H., & Howard, M. W. (2013). Optimally fuzzy temporal memory. *Journal of Machine Learning Research*, 14, 3753-3780.

Shankar, K. H., Singh, I., & Howard, M. W. (2015). Neural mechanism to simulate a scale-invariant future. *arXiv preprint arXiv:1503.03322*.

Tiganj, Z., Hasselmo, M. E., & Howard, M. W. (2015). A simple biophysically plausible model for long time constants in single neurons. *Hippocampus*, 25(1), 27-37.

References

Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., ... Burgess, N. (2006). The boundary vector cell model of place cell firing and spatial memory. *Reviews in Neuroscience*, 17(1-2), 71-97.

Chance, F. S., Abbott, L. F., & Reyes, A. D. (2002). Gain modulation from background synaptic input. *Neuron*, 35(4), 773-82.

Cole, R. P., Barnet, R. C., & Miller, R. R. (1995). Temporal encoding in trace conditioning. *Animal Learning & Behavior*, 23(2), 144-153.

Egorov, A. V., Hamam, B. N., Fransén, E., Hasselmo, M. E., & Alonso, A. A. (2002). Graded persistent activity in entorhinal cortex neurons. *Nature*, 420(6912), 173-8.

Egorov, A. V., Unsicker, K., & von Bohlen und Halbach, O. (2006). Muscarinic control of graded persistent activity in lateral amygdala neurons. *European Journal of Neuroscience*, 24(11), 3183-94.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1), 3-71.

Frank, L. M., Brown, E. N., & Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron*, 27(1), 169-178.

Fransén, E., Tahvildari, B., Egorov, A. V., Hasselmo, M. E., & Alonso, A. A. (2006). Mechanism of graded persistent cellular activity of entorhinal cortex layer V neurons. *Neuron*, 49(5), 735-46.

Hacker, M. J. (1980). Speed and accuracy of recency judgments for events in short-term memory. *Journal of Experimental Psychology: Human Learning and Memory*, 15, 846-858.

Hartley, T., Burgess, N., Lever, C., Cacucci, F., & O'Keefe, J. (2000). Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus*, 10(4), 369-79.

Howard, M. W., & Eichenbaum, H. (2013). The hippocampus, time, and memory across scales. *Journal of Experimental Psychology: General*, 142(4), 1211-30. doi: 10.1037/a0033621

Howard, M. W., & Eichenbaum, H. (2015). Time and space in the hippocampus. *Brain Research*, 1621, 345-354.

Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *Journal of Neuroscience*, 34(13), 4692-707. doi: 10.1523/JNEUROSCI.5808-12.2014

Howard, M. W., Shankar, K. H., Aue, W., & Criss, A. H. (2015). A distributed representation of internal time. *Psychological Review*, 122(1), 24-53.

Howard, M. W., Shankar, K. H., & Tiganj, Z. (2015). Efficient neural computation in the Laplace domain. In *Cognitive computation: Integrating neural and symbolic computation*.

Howard, M. W., Viskontas, I. V., Shankar, K. H., & Fried, I. (2012). Ensembles of human MTL neurons "jump back in time" in response to a repeated stimulus. *Hippocampus*, 22(9), 1833-1847.

Howard, M. W., Youker, T. E., & Venkatadass, V. (2008). The persistence of memory: Contiguity effects across several minutes. *Psychonomic Bulletin & Review*, 15(PMC2493616), 58-63.

Jung, M. W., Wiener, S. I., & McNaughton, B. L. (1994). Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *Journal of Neuroscience*, 14(12), 7347-56.

Kahana, M. J., Howard, M., & Polyn, S. (2008). Associative processes in episodic memory. In H. L. Roediger III (Ed.), *Cognitive psychology of memory, Vol. 2 of learning and memory - a comprehensive reference (J. Byrne, Editor)* (p. 476-490). Oxford: Elsevier.

Kjelstrup, K. B., Solstad, T., Brun, V. H., Hafting, T., Leutgeb, S., Witter, M. P., ... Moser, M. B. (2008). Finite scale of spatial representation in the hippocampus. *Science*, 321(5885), 140-3.

Komorowski, R. W., Garcia, C. G., Wilson, A., Hattori, S., Howard, M. W., & Eichenbaum, H. (2013). Ventral hippocampal neurons are shaped by experience to represent behaviorally relevant contexts. *Journal of Neuroscience*, 33(18), 8079-87. doi: 10.1523/JNEUROSCI.5458-12.2013

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *Journal of Neuroscience*, 29(31), 9771-7.

Longtin, A., Doiron, B., & Bulsara, A. (2002). Noise-induced divisive gain control in neuron models. *Biosystems*, 67.

Lubenov, E. V., & Siapas, A. G. (2009). Hippocampal theta oscillations are travelling waves. *Nature*, 459(7246), 534-9.

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, 71(4), 737-749.

Manning, J. R., Polyn, S. M., Litt, B., Baltuch, G., & Kahana, M. J. (2011). Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Science, USA*, 108(31), 12893-7.

Mehta, M. R., Lee, A. K., & Wilson, M. A. (2002). Role of experience and oscillations in transforming a rate code into a temporal code. *Nature*, 417(6890), 741-6.

Navaroli, V. L., Zhao, Y., Boguszewski, P., & Brown, T. H. (2011). Muscarinic receptor activation enables persistent firing in pyramidal neurons from superficial layers of dorsal perirhinal cortex. *Hippocampus*. doi: 10.1002/hipo.20975

O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381(6581), 425-428.

O'Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3), 317-30.

Patel, J., Fujisawa, S., Berényi, A., Royer, S., & Buzsáki, G. (2012). Traveling theta waves along the entire septotemporal axis of the hippocampus. *Neuron*, 75(3), 410-7. doi: 10.1016/j.neuron.2012.07.015

Post, E. (1930). Generalized differentiation. *Transactions of the American Mathematical Society*, 32, 723-781.

Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451), 585-90. doi: 10.1038/nature12160

Shankar, K. H. (2014). Quantum random walks and decision making. *Topics in cognitive science*, 6(1), 108–113.

Shankar, K. H. (2015). Generic construction of scale-invariantly coarse grained memory. *Lecture Notes in Artificial intelligence*, 8955, 175-184.

Shankar, K. H., & Howard, M. W. (2013). Optimally fuzzy temporal memory. *Journal of Machine Learning Research*, 14, 3753-3780.

Shankar, K. H., Singh, I., & Howard, M. W. (2015). Neural mechanism to simulate a scale-invariant future. *arXiv preprint arXiv:1503.03322*.

Silver, R. A. (2010). Neuronal arithmetic. *Nature Reviews Neuroscience*, 11(7), 474–489.

Tiganj, Z., Hasselmo, M. E., & Howard, M. W. (2015). A simple biophysically plausible model for long time constants in single neurons. *Hippocampus*, 25(1), 27-37.

van der Meer, M. A. A., & Redish, A. D. (2011). Theta phase precession in rat ventral striatum links place and reward information. *Journal of Neuroscience*, 31(8), 2843-54. doi: 10.1523/JNEUROSCI.4869-10.2011

1.

1. Report Type

Final Report

Primary Contact E-mail

Contact email if there is a problem with the report.

MarcWHoward777@gmail.com

Primary Contact Phone Number

Contact phone number if there is a problem with the report

315-256-4535

Organization / Institution name

Boston University

Grant/Contract Title

The full title of the funded effort.

A distributed representation of space and time

Grant/Contract Number

AFOSR assigned control number. It must begin with "FA9550" or "F49620" or "FA2386".

FA9550-12-1-0369

Principal Investigator Name

The full name of the principal investigator on the grant or contract.

Marc Howard

Program Manager

The AFOSR Program Manager currently assigned to the award

James Lawton

Reporting Period Start Date

07/01/2012

Reporting Period End Date

06/30/2015

Abstract

The goal of the award was to extend a mathematical framework for representing time to also construct scale-invariant representations of space. The extensive neurophysiological work done on the hippocampal place code (which resulted in a Nobel Prize in 2014) places physical constraints on this mechanism. We accomplished this goal by developing a computational framework that describes a wide range of functional cellular correlates in the hippocampus and related brain regions. Critically, this cellular-level model is connected to behavioral memory performance via cognitive models that take in the mathematical form of the representation. We have exceeded the initial goals of the project by developing a representation of numerosity as well as space using the same framework and a detailed model for function translation that is constrained by neurophysiological data from the hippocampal theta oscillation literature. This is an important development, as it opens up new frontiers in cognitive computation. Taken together, the work performed in this three year period was crucial in developing a mathematical model for cognitive operations that is constrained by neurophysiology. This model connects data from the subcellular level, to the systems neuroscience level to cognitive modeling. The mathematical framework is capable of forming the basis of a general brain-inspired cognitive computer.

Distribution Statement

This is block 12 on the SF288 form.
DISTRIBUTION A: Distribution approved for public release.

Distribution A - Approved for Public Release

Explanation for Distribution Statement

If this is not approved for public release, please provide a short explanation. E.g., contains proprietary information.

SF298 Form

Please attach your [SF298](#) form. A blank SF298 can be found [here](#). Please do not password protect or secure the PDF
The maximum file size for an SF298 is 50MB.

[AFD-070820-035.pdf](#)

Upload the Report Document. File must be a PDF. Please do not password protect or secure the PDF . The maximum file size for the Report Document is 50MB.

[SpaceGrantFinalReport.pdf](#)

Upload a Report Document, if any. The maximum file size for the Report Document is 50MB.

Archival Publications (published) during reporting period:

Howard, M. W., & Eichenbaum, H. (2013). The hippocampus, time, and memory across scales. *Journal of Experimental Psychology: General*, 142(4), 1211-30. doi: 10.1037/a0033621

Howard, M. W., & Eichenbaum, H. (2015). Time and space in the hippocampus. *Brain Research*, 1621, 345-354.

Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *Journal of Neuroscience*, 34(13), 4692-707. doi: 10.1523/JNEUROSCI.5808-12.2014

Howard, M. W., Shankar, K. H., Aue, W., & Criss, A. H. (2015). A distributed representation of internal time. *Psychological Review*, 122(1), 24-53.

Howard, M. W., Shankar, K. H., & Tiganj, Z. (2015). Efficient neural computation in the Laplace domain. In *Cognitive computation: Integrating neural and symbolic computation*.

Shankar, K. H. (2014). Quantum random walks and decision making. *Topics in cognitive science*, 6(1), 108–113.

Shankar, K. H. (2015). Generic construction of scale-invariantly coarse grained memory. *Lecture Notes in Artificial Intelligence*, 8955, 175-184.

Shankar, K. H., & Howard, M. W. (2013). Optimally fuzzy temporal memory. *Journal of Machine Learning Research*, 14, 3753-3780.

Shankar, K. H., Singh, I., & Howard, M. W. (2015). Neural mechanism to simulate a scale-invariant future. *arXiv preprint arXiv:1503.03322*.

Tiganj, Z., Hasselmo, M. E., & Howard, M. W. (2015). A simple biophysically plausible model for long time constants in single neurons. *Hippocampus*, 25(1), 27-37.

Changes in research objectives (if any):

We exceeded the initial goals of the application and developed a general framework for encoding scale-invariant (nonverbal) numerosity and a computational framework for general operations on these representations. This sets the stage for a general brain-inspired cognitive computer.

Change in AFOSR Program Manager, if any:

The initial program manager was Jay Myung. The current program manager is James Lawton.

Extensions granted or milestones slipped, if any:

N/A

AFOSR LRIR Number

LRIR Title

Reporting Period

Laboratory Task Manager

Program Officer

Research Objectives

Technical Summary**Funding Summary by Cost Category (by FY, \$K)**

	Starting FY	FY+1	FY+2
Salary			
Equipment/Facilities			
Supplies			
Total			

Report Document**Report Document - Text Analysis****Report Document - Text Analysis****Appendix Documents****2. Thank You****E-mail user**

Nov 09, 2015 01:52:15 Success: Email Sent to: MarcWHoward777@gmail.com